



Experiences of an IGTF Relying Party (FermiGrid – The Fermilab Campus Grid)

Keith Chadwick
Fermilab
chadwick@fnal.gov



What is FermiGrid?

FermiGrid is the Fermilab Campus Grid.

- We operate the central Fermilab Grid services cyberinfrastructure.
 - VOMRS, VOMS, GUMS, SAZ, Squid, Site Gatekeeper, etc.;
 - Services offered to CDF, D0, CMS and other Fermilab consumers;
 - The services are deployed in a Highly Available (HA) infrastructure with automatic failover (FermiGrid-HA).
- We operate multiple Grid resources for a variety of client communities.
 - CDF, CMS & D0 experiments;
 - Other smaller Fermilab experiments.
- We coordinate and interoperate with other campus Grids (Purdue), regional Grids (NYSGrid, SuraGrid), and national cyberinfrastructures (OSG, TeraGrid).
- <http://fermigrid.fnal.gov>



Inventory of Physical Hardware, Virtual Systems and Services

	Physical Systems	Virtual Systems	Virtualization Technology	Service Count
FermiGrid-HA Services	6	34	Xen	17
CDF, D0, GP Gatekeepers	9	28	Xen	9+6
Fermi & OSG Gratia	4	10	Xen	12
OSG ReSS	2	8	Xen	2
Integration Test Bed (ITB)	2+8	14+32	Xen	14
Grid "Access" Services	2	4	Xen	4
"FermiCloud"	8 (+16)	64 (+128)	Xen	--
"Fgtest" Systems	7	51	Xen	varies
"Cdf Sleeper Pool"	3	9	Xen	1+1
"GridWorks"	11	~20	Kvm	1



"Owned" Job Slots by Client Community

Community	# Clusters	# Gatekeeper	# Slots	VO Location
_____	_____	_____	_____	_____
CDF Experiment	3	5*	5,315	Fermilab
CMS Experiment	1	4*	5,144	CERN
D0 Experiment	2	2	5,597	Fermilab
"Other" Fermilab	1	3	965	Fermilab
_____	_____	_____	_____	_____
Total	7	14	17,021	

Table Data Source = <http://fermigrid.fnal.gov/fermigrid-metrics.html>

- Four of the five CDF gatekeepers are open for opportunistic use.
- Only one of the four CMS gatekeepers is open for opportunistic use.



Use of IGTF Infrastructure

FermiGrid operates a large number of systems that are “relying partners” of the IGTF infrastructure.

We configure our systems that support “HA” services to update CRL’s every hour.

We configure our clusters to use a central CRL update mechanism that updates CRL’s every hour.

We have a centrally managed high availability squid service (squid-ha) to cache CRL updates.

But, most Certificate Authorities are not publishing their CRL lifetimes in a manner that is “squid friendly”.

Consequently, we have had to establish a set of cache refresh parameters in the squid-ha servers.




Default CRL Cache Lifetimes



We have established the following “default” CRL cache lifetimes when the CA does not specify them:

```
# TAG: refresh_pattern
#      usage: refresh_pattern [-i] regex min percent max [options]
#      The refresh_pattern lines are checked in the order listed here.
refresh_pattern ^ftp:                1440      20%      10080
refresh_pattern ^gopher:             1440      0%       1440
refresh_pattern \.crl$                5         25%       120
refresh_pattern \.der$                5         25%       120
refresh_pattern \.pem$                5         25%       120
refresh_pattern \.r0$                 5         25%       120
refresh_pattern \.pacman$             5         10%      1440
refresh_pattern .                     5         20%      4320
```


Squid-HA - Calls per Day




Squid-HA - Clients per Day




Squid-HA - CRL Downloads Hit/Miss Rate



Squid-HA Status - CRL Downloads



Squid-HA Status - CRL Downloads





Squid Cache Statistics & Benefits

Average CRL downloads through squid server / day	2,121,174
Average squid cache CRL hit rate (past month)	95%
CRLs served from squid cache	2,018,970
CRL's actually downloaded	102,204
Average # of CRLs per CA downloaded by FermiGrid / day	1,087

The FermiGrid Squid cache benefits FermiGrid as well as *all* of the IGTF CA's

CRL x509 Last Update Time:

- `openssl crl ${hash}.r0 | grep 'Last Update'`

CRL x509 Next Update Time:

- `openssl crl ${hash}.r0 | grep 'Next Update'`
-

CRL Web Modification Time:

- `curl -D a.a $crl_url &> /dev/null ; grep -a 'Modified' a.a`

CRL Web Expiration Time:

- `curl -D a.a $crl_url &> /dev/null ; grep -a 'Expires' a.a`

CRL Web Cache Lifetime:

- `curl -D a.a $crl_url &> /dev/null ; grep -a 'max-age' a.a`



The Results of a Survey of IGTF CAs

Survey of IGTF Accredited CAs at 1254868843 (Tue Oct 6 17:40:43 CDT 2009)	
Number of CA's	94
Number of CA's that failed CRL download	0
Number of CRL's with openssl Last Update Times	94
Number of CRL's with openssl Next Update Times	94
<hr/>	
Number of CRL's with Web Modification Times in http header	81
Number of CRL's with Web Expiration Times in http header	11
Number of CRL's with explicit cache lifetime (max-age) in the http header	13
Number of CRL's without Web Modification, Expiration Time or cache lifetime	13



So – What are the issues? – #1

Most Certificate Authorities are not publishing their CRL lifetimes in a manner that is “squid friendly”;

- FermiGrid has to “guess” at appropriate default values for the CRL cache lifetime [120 minutes = 7,200 seconds];
- The 13 CA’s that are publishing max-age cache lifetimes in the http headers are using lifetimes of either 3,600 or 86,400 seconds.



Example of a "Good" CRL Publication

```
$ curl -D a.a `cat /etc/grid-security/certificates/28a58577.crl_url` &> /dev/null
$ cat a.a
HTTP/1.1 200 OK
Date: Tue, 06 Oct 2009 23:09:09 GMT
Server: Apache/2.2.3 (FreeBSD) mod_ssl/2.2.3 OpenSSL/0.9.7e-p1 PHP/5.2.0 with
       Suhosin-Patch DAV/2 SVN/1.4.2 Phusion_Passenger/2.0.6
Last-Modified: Thu, 24 Sep 2009 15:05:32 GMT
ETag: "20b574-1b5-29ab2700"
Accept-Ranges: bytes
Cache-Control: max-age=86400
Expires: Wed, 07 Oct 2009 23:09:09 GMT
Content-Type: text/plain
Content-Length: 437
Connection: Keep-Alive
Age: 0
```


Google search:

- "set max-age http header using apache"

Top hit:

- <http://www.askapache.com/htaccess/apache-speed-cache-control.html>

Other relevant hits:

- http://www.mnot.net/cache_docs/#EXPIRES
- http://www.mnot.net/cache_docs/#CACHE-CONTROL



So – What are the issues? – #2

When a Certificate Authority is unavailable;

- The FermiGrid system administrators receive ~3,000 email messages per day about CRL update failures;
- Other “real” incidents can be lost or overlooked in this deluge of error messages.



CRL Download Incident 1

"Dear <name>,"

We are aware of the situation. Last Wednesday, all grid services were shut down in response to severe security incident at the Institute. All hardware and software resources are currently being audited. I will talk about this in more details at wed conference call.

The CRLs however were relocated to vetted resources and were down for around 30 hours until Thursday night. However it appears that they may have continued to be unreachable over the weekend. While we could access the CRL we had noticed the nagios monitor couldn't and initially thought this was due to DNS caching delays. Since then we have suffered by severed fibre cables and power outages. When things go bad they real bad fast. We have been and are still investigating the problem but we expect to have normal service resumed by tomorrow (tuesday), fingers crossed."

CRL Download Incident 2

“One of the CA operators revoked a certificate on a test CA running on the same machine as the production CA, and the test CA's CRL overwrote the production CA's CRL, thereby causing the CRL errors.”



CRL Download Incident 3

On Friday 25-Sep-2009, the issuer of the 295adc19.r0 (REUNA-ca)

Issued a CRL with a malformed last update time. My system logs reported:

```
fetch-crl[11909]: 20090925T163505-0500 Warning: CRL downloaded from has lastUpdate time in the future. Verify local clock and inspect 295adc19.r0.  
fetch-crl[11909]: 20090925T163505-0500 CRL 295adc19.r0 replaced with downloaded one, since current one has a lastUpdate time in the future.
```

```
$ openssl crl -in /etc/grid-security/certificates/295adc19.r0 -text | head -n 8
```

Certificate Revocation List (CRL):

Version 2 (0x1)

Signature Algorithm: sha1WithRSAEncryption

Issuer: /C=CL/O=REUNACA/CN=REUNA Certification Authority

Last Update: Sep 25 23:20:02 2009 GMT

Next Update: Oct 25 23:20:02 2009 GMT

CRL extensions:

X509v3 Authority Key Identifier:

keyid:A3:AE:48:8E:B9:C1:8E:B1:92:AA:5E:0C:D0:DC:9D:4B:05:2E:2C:57

The date in the CRL is:

```
$ date -d "Sep 25 23:20:02 2009 GMT"
```

Fri Sep 25 18:20:02 CDT 2009

The current date is:

```
$ date
```

Fri Sep 25 16:58:40 CDT 2009

CRL Download Incident 4

fetch-crl[21853]: 20091007T094013-0500

- RetrieveFileByURL: download no data from <http://gridca.ihep.ac.cn/cacrl.pem>

fetch-crl[21853]: 20091007T094013-0500

- Persistent errors (219 hours) for ba2f39ca:

fetch-crl[21853]: 20091007T094013-0500

- Could not download any CRL from /usr/local/vdt-tomcat/globus/TRUSTED_CA//ba2f39ca.crl_url:

fetch-crl[21853]: 20091007T094013-0500

- download failed from '<http://gridca.ihep.ac.cn/cacrl.pem>'



CRL Download Incident 5

On 14-Jul-2009, the Squid logs (Squid 2.6STABLE18) started filling up with the following messages:

```
2009/07/14 14:11:19| httpReadReply: Excess data from "GET http://ca.ncsa.uiuc.edu/e8ac4b61.crl"
2009/07/14 14:11:19| httpReadReply: Excess data from "GET http://ca.ncsa.uiuc.edu/e8ac4b61.crl"
```

There is a difference between a wget and a curl for this file. The wget is getting an extra byte added. Squid 2.6STABLE18 (installed from the VDT cache) is detecting this.

```
-$ wget -O e8ac4b61.crl wget http://ca.ncsa.uiuc.edu/e8ac4b61.crl
--15:37:52-- http://ca.ncsa.uiuc.edu/e8ac4b61.crl
      => `e8ac4b61.crl.wget'
Resolving ca.ncsa.uiuc.edu... 141.142.15.53
Connecting to ca.ncsa.uiuc.edu[141.142.15.53]:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 476 [application/pkix-crl]
```

```
100%
[=====] 477
```

```
15:37:52 (4.55 MB/s) - `e8ac4b61.crl.wget' saved [477/476]
```

```
$ curl http://ca.ncsa.uiuc.edu/e8ac4b61.crl > e8ac4b61.crl.curl
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload   Total   Spent    Left  Speed
100  476  100  476    0     0  4544    0 --:--:-- --:--:-- --:--:--    0
```

```
$ ls -acl e8ac4b61.crl*
-rw-r--r-- 1 chadwick chadwick 476 Jul 14 15:39 e8ac4b61.crl.curl
-rw-r--r-- 1 chadwick chadwick 477 Jul 14 15:37 e8ac4b61.crl.wget
```

16-Oct-2009

FermiGrid - Tagpma

22



The FermiGrid Request to CA Operators

Please:

1. Add the appropriate http headers to specify your CRL modification time, expiration time and maximum cache age.
 - Don't specify "no-cache" on your http header.
2. Don't shut your CA down without establishing an "alternate" location for the CRL downloads;
 - Especially when you may be / are having a security incident!
3. Verify the changes to your CA infrastructure;
 - Especially immediately after publishing new CRLs.
4. Monitor your CA infrastructure overnight and the weekend.
5. Have a disaster recovery plan;
 - And test it periodically!

Any questions?